Testimony and Statement for the Record of
Bruce Schneier

Chief Technical Officer, Counterpane Internet Security, Inc.

Hearing on
Internet Security

Before the

Subcommittee on Science, Technology, and Space of the
Committee on Commerce, Science and Transportation

United States Senate

July 16, 2001
253 Russell Senate Office Building

My name is Bruce Schneier. I am the founder and Chief Technical Officer of Counterpane Internet Security. Inc. Counterpane was founded to address the immediate need for increased Internet security, and essentially provides burglar alarm services for computer networks. I am the author of seven books on cryptography and computer security, as well as hundreds of articles and papers on those topics. For several years, I have been a security consultant to many major Internet companies.

I'd like to thank the Committee for holding this hearing today. Internet security is an enormously important issue, and one that will become increasingly important as the Internet affects the lives of more people. Simply stated, during the last decade the Internet has transitioned from a technological plaything for a few people to a critical infrastructure as fundamental as the phone system. Internet security has transitioned from an academic curiosity to a fundamental enabling technology for our future. The limits of Internet security are the limits of the Internet, and the limits of the Internet profoundly affect our country as the Information Economy continues to grow.

I believe that there are two questions before the Committee today. The first is whether the Internet is safe enough to conduct business on. The second, if you agree that the Internet is not safe enough, is what we can do to improve the situation. I will focus my remarks on these two issues.

Introduction

The Internet is critical to business. Companies have no choice but to connect their internal networks to the rest of the world—to link with customers, suppliers, partners, and their own employees. But with that connection comes new threats: malicious hackers, criminals, industrial

spies. These network predators regularly steal corporate assets and intellectual property, cause service breaks and system failures, sully corporate brands, and frighten customers. Unless companies can successfully navigate around these, they will not be able to unlock the full business potential of the Internet.

Traditional approaches to computer security center around preventive techniques, and they don't work. Despite decades of research, and hundreds of available security products, the Internet has steadily become more dangerous. The increased complexity of the Internet and its applications, the rush to put more services and people on the Internet, and the desire to interconnect everything all contribute to the increased insecurity of the digital world.

Security based solely on preventive products is inherently fragile. Newly discovered attacks, the proliferation of attack tools, and flaws in the products themselves all result in a network becoming vulnerable at random (and increasingly frequent) intervals.

Active security monitoring is a key component missing in most networks. Insurance is another. In business, insurance is the risk manager of last resort. And in most cases, insurance drives security requirements. Companies install a burglar alarm system in their warehouse not because it reduces theft, but because it reduces their insurance rates. As the need for Internet security becomes more universally recognized , insurance companies will begin to drive security requirements and demand product improvements.

The third key component to a secure Internet is law enforcement. The primary reason we live in a safe society is that we prosecute criminals. Today the Internet is a lawless society; hackers can break into computers with relative impunity. We need to turn the Internet into a lawful society, through regular prosecution and conviction of Internet criminals.

The Importance of Security

When I began working in computer security, the only interest was from the military and a few scattered privacy advocates. The Internet has changed all that. The promise of the Internet is to be a mirror of society. Everything we do in the real world, we want to do on the Internet: conduct private conversations, keep personal papers, sign letters and contracts, speak anonymously, rely on the integrity of information, gamble, vote, publish digital documents. All of these things require security. Computer security is a fundamental enabling technology of the Internet; it's what transforms the Internet from an academic curiosity into a serious business tool. The limits of security are the limits of the Internet. And no business or person is without these security needs.

The risks are real. Everyone talks about the direct risks: theft of trade secrets, customer information, money. People also talk about the productivity losses due to computer security problems. What's the loss to a company if its e-mail goes down for two days? Or if ten people have to scramble to clean up after a particularly nasty intrusion? I've seen figures as high as $10 billion quoted for worldwide losses due to the ILOVEYOU virus; most of that is due to these productivity losses.

More important are the indirect risks: loss of customers, damage to brand, loss of goodwill. Last year Egghead.com had a network break-in and it was rumored that a million credit card numbers were stolen. Regardless of how the investigation turned out, some percentage of customers decided to shop elsewhere. When CD Universe suffered a credit card theft in early 2000, it cost them dearly in their war for market share against Amazon.com and CDNow. In the aftermath of the Microsoft attack in October 2000, the company spent much more money and effort containing the public relations problem than fixing the security problem. The public perception that their source code was untainted was much more important than any effects of the actual attack.

And more indirect risks are coming. European countries have strict privacy laws; American companies can be held liable if they do not take steps to protect the privacy of their European customers. While "safe harbor" provisions may provide immediate relief, it will not solve the problem once the European countries realize that their data is not being protected.

The U.S. has similar laws in particular industries—banking and healthcare—and there are bills in Congress to protect privacy more generally. We have not yet seen shareholder lawsuits against companies that failed to adequately secure their networks and suffered the consequences, but they're coming. Can company officers be held personally liable if they fail to provide for network security? The courts will be deciding this question in the next few years.

As risky as the Internet is, companies have no choice but to be there. The lures of new markets, new customers, new revenue sources, and new business models are just so great that companies will flock to the Internet regardless of the risks. There is no alternative. This, more than anything else, is why computer security is so important.

The Failure of Traditional Security

Five years ago, network security was relatively simple. No one had heard of denial-of-service attacks shutting down Web servers, Web page scripting flaws, or the latest vulnerabilities in Microsoft Outlook Express. In recent years came intrusion detection systems, public-key infrastructure, smart cards, VPNs, and biometrics. New networking services, wireless devices, and the latest products regularly turn network security upside down. There are literally hundreds of network security products you can buy, and they all claim to provide you with security. They regularly fail, but still you hear companies say: "Of course I'm secure. I bought a firewall."

Network security is an arms race, and the attackers have all the advantages. First, network defenders occupy what military strategists call "the position of the interior": the defender has to defend against every possible attack, while the attacker only has to find one weakness. Second, the immense complexity of modern networks makes them impossible to properly secure. And third, skilled attackers can encapsulate their attacks in software, allowing people with no skill to use them. It's no wonder businesses can't keep up with the threat.

What's amazing is that no one else can either. Computer security is a 40-year-old discipline; every year there's new research, new technologies, new products, even new laws. And every year things get worse.

If there's anything computer security professionals have learned about the Internet, it's that security is relative. Nothing is foolproof. What's secure today may be insecure tomorrow. Even companies like Microsoft can get hacked, badly. There are no silver bullets. The way forward is not more products, but better processes. We have to stop looking for the magic preventive technology that will avoid the threats, and embrace processes that will help us manage the risks.

Security and Risk Management

Ask any network administrator what he needs security for, and he can describe the threats: Web site defacements, corruption and loss of data due to network penetrations, denial-of-service attacks, viruses and Trojans. The list seems endless, and the endless slew of news stories prove that the threats are real.

Ask that same network administrator how security technologies help, and he'll discuss avoiding the threats. This is the traditional paradigm of computer security, born out of a computer science mentality: figure out what the threats are, and build technologies to avoid them. The conceit is that technologies can somehow "solve" computer security, and the end result is a security program that becomes an expense and a barrier to business. How many times has the security officer said: "You can't do that; it would be insecure"?

This paradigm is wrong. Security is a people problem, not a technology problem. There is no computer security product-or even a suite of products-that acts as magical security dust, imbuing a network with the property of "secure." It can't be done. And it's not the way business works.

Businesses manage risks. They manage all sorts of risks; network security is just another one. And there are many different ways to manage risks. The ones you choose in a particular situation depend on the details of that situation. And failures happen regularly; many businesses manage their risks improperly, pay for their mistakes, and then soldier on. Businesses are remarkably resilient.

To take a concrete example, consider a physical store and the risk of shoplifting. Most grocery stores accept the risk as a cost of doing business. Clothing stores might put tags on all their garments and sensors at the doorways; they mitigate the risk with a technology. A jewelry store might mitigate the risk through procedures: all merchandise stays locked up, customers are not allowed to handle anything unattended, etc. And that same jewelry store will carry theft insurance, another risk management tool.

More security isn't always better. You could improve the security of a bank by strip-searching everyone who walks through the front door. But if you did this, you would have no business. Studies show that most shoplifting at department stores occurs in dressing rooms. You could improve security by removing the dressing rooms, but the losses in sales would more than make

up for the decrease in shoplifting. What all of these businesses are looking for is adequate security at a reasonable cost. This is what we need on the Internet as well—security that allows a company to offer new services, to expand into new markets, and to attract and retain new customers. And the particular computer security solutions they choose depend on who they are and what they are doing.

Detection and Response

Most computer security is sold as a prophylactic: encryption prevents eavesdropping, firewalls prevent unauthorized network access, PKI prevents impersonation. To the world at large, this is a strange marketing strategy. A door lock is never sold with the slogan: “This lock prevents burglaries.” No one ever asks to purchase “a device that will prevent murder.” But computer security products are sold that way all the time. Companies regularly try to buy “a device that prevents hacking.” This is no more possible than an anti-murder device.

When you buy a safe, it comes with a rating. 30TL—30 minutes, tools. 60TRTL—60 minutes, torch and tools. What this means is that a professional safecracker, with safecracking tools and an oxyacetylene torch, can break open the safe in an hour. If an alarm doesn’t sound and guards don’t come running within that hour, the safe is worthless. The safe buys you time; you have to spend it wisely.

Real-world security includes prevention, detection, and response. If the prevention mechanisms were perfect, you wouldn’t need detection and response. But no prevention mechanism is perfect. This is especially true for computer networks. All software products have security bugs, most network devices are misconfigured, and users make all sorts of mistakes. Without detection and response, the prevention mechanisms only have limited value. They’re fragile. And detection and response are not only more cost effective, but also more effective, than piling on more prevention.

On the Internet, this translates to monitoring. In October 2000, Microsoft discovered that an attacker had penetrated their corporate network weeks before, and might have viewed or even altered the source code for some of their products. Administrators discovered this breach when they noticed twenty new accounts being created on a server. Then they went back through their network’s audit logs and pieced together how the attacker got in and what he did. If someone had been monitoring those audit logs—automatically generated by the firewalls, servers, routers, etc.—in real time, the attacker could have been detected and repelled at the point of entry.

That’s real security. It doesn’t matter how the attacker gets in, or what he is doing. If there are enough motion sensors, electric eyes, and pressure plates in your house, you’ll catch the burglar regardless of how he got in. If you are monitoring your network carefully enough, you’ll catch a hacker regardless of what vulnerability he exploited to gain access. And if you can respond quickly and effectively, you can repel the attacker before he does any damage. Good detection and response can make up for imperfect prevention.

And real security is about people. On the day you're attacked, it doesn't matter how your network is configured, what kind of boxes you have, or how many security devices you've installed. What matters is who is defending you.

Prevention systems are never perfect. No bank ever says: "Our safe is so good, we don't need an alarm system." No museum ever says: "Our door and window locks are so good, we don't need night watchmen." Detection and response are how we get security in the real world, and they're the only way we can possibly get security on the Internet. We must invest in network monitoring if we are to properly manage the risks associated with our nation's network infrastructure.

Insurance

Eventually, the insurance industry will subsume the computer security industry. Not that insurance companies will start marketing security products, but rather that the kind of firewall you use—along with the kind of authentication scheme you use, the kind of operating system you use, and the kind of network monitoring scheme you use—will be strongly influenced by the constraints of insurance.

Consider security, and safety, in the real world. Businesses don't install building alarms because it makes them feel safer; they do it because they get a reduction in their insurance rates. Building owners don't install sprinkler systems out of affection for their tenants, but because building codes and insurance policies demand it. Deciding what kind of theft and fire prevention equipment to install are risk management decisions.

The risk taker of last resort is the insurance industry, and businesses achieve security through insurance. They take the risks they are not willing to accept themselves, bundle them up, and pay someone else to make them go away. If a warehouse is insured properly, the owner is significantly less worried about fire or other disasters. Similarly, if a network is insured properly, the owner is significantly less worried about the hacking risks.

This is the future. Concerned about denial-of-service attacks? Get bandwidth interruption insurance. Concerned about data corruption? Get data integrity insurance. (I'm making these policy names up, here.) Concerned about negative publicity due to a widely publicized network attack? Get a rider on your good name insurance that covers that sort of event. The insurance industry isn't offering all of these policies yet, but it is coming.

The effects of this change will be considerable. Every business will have network security insurance, just as every business has insurance against fire, theft, and any other reasonable threat. To do otherwise would be to behave recklessly and be open to lawsuits. Details of network security become check boxes when it comes time to calculate the premium. Do you have a firewall? Which brand? Your rate may be one price if you have this brand, and a different price if you have another brand. Do you have a service monitoring your network? If you do, your rate goes down this much.

This process changes everything. What will happen when the CFO looks at his premium and realizes that it will go down 50% if he gets rid of all his insecure Windows operating systems and replaces them with a secure version of Linux? The choice of which operating system to use will no longer be 100% technical. Microsoft, and other companies with shoddy security, will start losing sales because companies don't want to pay the insurance premiums. In this vision of the future, how secure a product is becomes a real, measurable, feature that companies are willing to pay for...because it saves them money in the long run. Already some insurance companies are starting to do this.

Other systems will be affected, too. Online merchants and brick-and-mortar merchants will have different insurance premiums, because the risks are different. Businesses can add authentication mechanisms—public-key certificates, biometrics, smart cards—and either save or lose money depending on their effectiveness. Computer security "snake-oil" peddlers who make outlandish claims and sell ridiculous products will find no buyers as long as the insurance industry doesn't recognize their value. In fact, the whole point of buying a security product or hiring a security service will not be based on threat avoidance; it will be based on risk management.

And it will be about time. Sooner or later, the insurance industry will sell everyone anti-hacking policies. It will be unthinkable not to have one. And then we'll start seeing good security rewarded in the marketplace.

Law Enforcement

The primary reason we feel safe walking the streets of our country is because criminals are arrested and prosecuted. In areas where prosecution is less common, the streets are more dangerous. In countries where prosecution is rare or arbitrary, criminals run rampant. This same thinking must be applied to the Internet.

Right now, most criminal hackers can operate with impunity, and they know that. Most Internet crimes are never discovered by the victims. Of those that are known, most are covered up. Of those that are made public, most never result in arrests, let alone convictions. The Internet is still a lawless environment.

This needs to change. Prosecution and conviction of criminals has two effects. One, it sends a clear message to everyone else. And two, it takes the convicted criminals out of circulation during their incarceration. Both of these things act as a deterrence.

One of the best things that happened for Internet security in the year 2000 was the series of high-profile prosecutions and convictions. This has had a visible chilling effect on some hacking groups. But more is required.

This is not easy. The Internet was not designed to aid forensic analysis, and many types of hacks are not currently traceable. Jurisdiction is also a problem; our criminal justice system is not designed to deal with criminals who can be anywhere in the world while attacking someone in another part of the world. But we need to do it.

Conclusion

Network security risks will always be with us. The downside of being in a highly connected network is that we are all connected with the best and worst of society. Security products will not solve the problems of Internet security, any more than they solve the security problems in the real world. The best we can do is to manage the risks: employ technological and procedural mitigation while at the same time allowing businesses to thrive.

Security equals vigilance, a day-to-day process. There are hundreds of technological solutions, but none that will ultimately fix the problem. It's been thousands of years, and the world still isn't a safe place. There is no way to "solve" the burglary problem. There is no device you can buy to prevent murder. No matter how fast technology advances, guards and alarms are still state-of-the-art.

The key to effective security is human intervention. Automatic security is necessarily flawed. Smart attackers bypass the security, and new attacks fool products. People are needed to recognize, and respond to, new attacks and new threats. It's a simple matter of regaining a balance of power: human minds are the attackers, so human minds need to be the defenders as well.

I believe that the Internet will never be totally secure. In fact, I believe that the Internet will continue to get less and less secure as it gets more interesting, more useful, and more valuable. Just like the real world, security is a process. And the processes of detection and response, risk management and insurance, and forensics and prosecution will serve the Internet world just as they serve the real world.